I hereby certify that this paper is being deposited with the United States Postal Service as Express Mail in an envelope addressed to: Assistant Commissioner for Patents, Washington, D.C. 20231, on this date.

10-14-98
Date

Express Mail Label No. EM058839456US

96RSS017,075/73063

NEURAL NETWORK FOR ~~CONTROLLING~~ CALLS IN A TELEPHONE SWITCH

Field of the Invention

The field of the invention relates to telephony systems and more particularly to automatic call distributors used with private networks.

Background of the Invention

Automatic call distribution systems are known. Such systems are typically used in an organizational context as a means of distributing telephone calls among a group of agents of the organization.

Often the organization disseminates a single telephone number to its customers and to the public in general as a means of contacting the organization. As calls are directed to the organization from the public switch telephone network (PSTN), the automatic call distribution system directs the calls to its agents based upon some algorithm, typically based upon availability. For example, where all agents are consider equal, the automatic call distributor (ACD) may distribute the calls based upon which agent position (telephone) has been idle the longest.

In order to distribute incoming calls from the PSTN to the available agents, the interaction of a controlling computer with a switching fabric of the ACD becomes essential. Often a connection of the ACD to a local PSTN is in the form of a number of trunk connections. Each of the trunk connections is monitored by the controller for incoming calls. Where a call is detected, the controller searches for and

selects an idle agent. Upon selecting an agent, the controller instructs the switch to form a connection between the incoming trunk and selected agent.

In more complicated systems, the organization may use a number of telephone numbers to identify different individuals and functions within the organization. Each telephone number may be assigned to a particular incoming trunk or group of incoming trunk lines. As such, the controller may be required to recognize a call target based upon an identity of an incoming trunk line and route the call accordingly.

In other systems, the ACD of an organization may receive calls directed to different call targets over the same trunk lines. In such a case, the call target may be identified to the ACD by a pulse code modulated (PCM) signal transferred from the PSTN to the controller of the ACD by a dialed number identification service (DNIS) operating from within the PSTN.

In systems associated with service organizations, where many calls are received and handled by many agents, it may be important for an agent to have ready access to customer files. In such a situation, a database is maintained of existing customers. Customer records may be displayed on agent terminals as the agents converse with specific customers. In some cases, the customer may be identified to the database for display of records on the terminal by the agent entering a customer identifier into a keyboard associated with the terminal. Alternatively, the controller of the ACD may transfer an identifier of the customer to the database based upon an automatic number

identification (ANI) facility, operating from within the PSTN.

Where ANI is used, the controller of the ACD receives the ANI digits (identifying the caller via the caller's telephone number) at the same time the call arrives from the PSTN. Upon selecting an agent, the controller may transfer a call to a queue of the selected agent (or group of agents) or directly to the selected agent. At the same time that the call is delivered to the agent, the controller sends an identifier of the selected agent and ANI number of the customer to a controller of the database (the host). The host, in turn, displays the customer records on a computer monitor of the terminal of the selected agent at the same time the call is delivered.

While the existing method of ACD operation is relatively satisfactory, the relative loading of agent groups varies widely. Often sales campaigns are launched involving some agents of a selected number of agent groups. Other times, agents may be drafted for outbound call campaigns based upon availability. However, the varying levels of call processing may intermittently leave agents handling incoming calls idle. Other times, incoming calls may overwhelm agents designated for accepting such calls. Accordingly, a need exists for a means of dynamically adapting operating resources of automatic call distributors to changing circumstances, such that the dynamic nature of call processing may be accommodated without difficulty.

Summary

A method and apparatus are provided for processing calls in an automatic call distributor. The method includes the steps of learning a set of desired resource relationships for servicing a plurality of call processing load conditions in the automatic call distributor and, afterwards, distributing resources of the automatic call distributor based upon call processor loading and the learned desired set of resource relationships.

Brief Description of the Drawings

FIG. 1 is a block diagram of an automatic call distributor in accordance with an illustrated embodiment of the invention;

FIG. 2 is a flow chart of process steps performed by the system of FIG. 1;

FIG. 3 is a neural network of the system of FIG. 1; and

FIG. 4 is a neural network of the system of FIG. 1.

Detailed Description of a Preferred Embodiment

FIG. 1 is a block diagram of an automatic call distributor (ACD) system 10, generally, in accordance with an illustrated embodiment of the invention. The ACD 10 may include a call distributor (CD) 20 coupled to the PSTN 14 through a number of trunk lines 30. The CD 20 may route calls received from the PSTN 14 to a number of agent consoles 28 (one shown in FIG. 1) under the control of a host 16. The host 16 may include one or more neural networks (NN) 32. Each agent console 28

has associated with it, an agent terminal 18 coupled to a database within the host 16.

Under the embodiment, the ACD 10 may direct calls from a number of customers 12 (one shown in FIG. 1) to a number of agents (one agent station 18, 28 shown in FIG. 1), or calls from a number of agents to a number of customers under the control of the NN 32.

Calls from customers 12 may be initiated from subscriber locations anywhere within the PSTN 14 and routed by the PSTN 14 to the ACD 10. To facilitate delivery of calls, an owner of the ACD 10 may disseminate by advertising or otherwise, an address for delivery of calls. In the case of calls, the address may simply be a telephone number.

Under the illustrated embodiment, calls delivered from the PSTN 14 to the ACD 10 under voice format may be handled conventionally. The ACD 10 may include a switch 26, which may be interconnected with the PSTN 14 through a number of trunk lines 30. The PSTN 14 may offer service on the trunk lines 30 in association with services such as ANI or DNIS. Call control, call maintenance, and call set-up may be accomplished over the trunk line itself or over an associated control channel.

DNIS information supplied by the PSTN 14 is useful where inbound calls to the ACD 10 may be directed to any of a large block of telephone numbers assigned to the ACD 10. Calls of the block of numbers may be delivered to the ACD 10 through the trunk lines 30 in rotary fashion, so that when the calling party from the PSTN appears, for example, on trunk T1, it can be determined whether the calling party was, in fact,

calling the telephone number corresponding to trunk T1 or was, in fact, calling the telephone number corresponding to trunk T2 and was rotated down to the next available trunk, T1.

The switch 26 is controlled by a central processing unit, or CPU 22, in conjunction with peripheral memory device 24. Control of the switch 26 and communications with the host 16 and PSTN 14 may be accomplished generally as described in U.S. Patent No. 5,268,903, and U.S. Patent No. 5,140,611, both to Jones, and both incorporated herein by reference. Routing of calls to agents 28 and overflow of calls may be accomplished generally as described in: U.S. Patent No. 5,335,269 to Steinlicht et al.; U.S. Patent No. 5,365,581 to Baker et al.; and U.S. Patent No. 5,384,841 to Adams et al., all incorporated herein by reference.

During operation, the CPU 22 monitors each port of the switch 26 for changes in status. A change in status may be an agent unit 28 going off-hook to make a call, an agent unit hanging up after a call, or it may be a call alerting tone detected on a trunk, alerting the CPU 22 to the presence of an incoming call.

Where the status change is an agent 28 hanging up, the CPU 22 acts to tear-down the call connection within the switch 26 between the agent at a first port of the switch and a second party to the conversation communicating through a second port of the switch 26. Upon tear down of the connection, the CPU 22 also sends a message to the host 16, notifying the host of termination of the call connection. The message to the host 16 includes at least the identity of the agent 28.

Where the status change is a call alert signal on an incoming trunk line (or control channel associated with the incoming trunk line), the CPU 22 may send a call arrival message to the host 16 (and NN 32). Based upon system loading, the host 16 may authorized the CD 20 to accept the call. Upon receiving authorization to accept the call, the CD 20 sends an acknowledge message to the PSTN 14 accepting the call. The PSTN 14 may respond by forwarding DNIS and ANI information, identifying the called and calling party.

Upon accepting the call, the CPU 22 first stores the DNIS and ANI numbers in a termination table of the memory 24. More specifically, the CPU 22 maintains a table of call information for each port of the switch 26. Where a call is accepted on an incoming trunk line, the CPU 22 enters the DNIS and ANI number into the table for the incoming trunk line upon which the call is received.

In addition to updating the termination table within memory 24, the CPU 22 also generates a call identifier (also sometimes referred to as a call ID or sequence number) for the call, unique to the switch 26. The call identifier along with the ANI and DNIS numbers may then be sent to the host 16 as part of a call arrival message. The call arrival message includes a header identifying the message as being a call arrival message. The message may also include an ANI number, DNIS information and the call identifier.

Delivery of the ANI and DNIS numbers and call identifier allows the host 16 to create a unique call record for the call in memory 24, in a call record area of memory 24. The call record (and specifically the

ANI information within the call record) may be used to retrieve customer records for delivery to an appropriate display terminal 18 once the call has been assigned to an agent 28.

The CPU 22 then, by reference to the DNIS number, determines the identity of an agent 28 to which the call is to be directed. For example, the DNIS number may be used to differentiate between calls directed to a first telephone number arriving on a first incoming trunk group directed to a sales group of the organization from calls directed to a service group of the organization. Since agents servicing sales calls would, in most cases, not handle calls directed to service, the DNIS number provides a convenient means of differentiating between two or more types of calls.

The CPU 22 may maintain a series of lookup tables within memory 24 identifying groups and agents within each group. A header of each lookup table may identify the DNIS numbers of calls to be directed to the group.

Each group lookup table may contain a list of identifiers of agents assigned to the group. Agents may be added to or deleted from a group under control of the NN 32.

Also contained within each group list is a status indicator of each agent. Agents may be selected based upon availability or qualifications.

Upon determining the identity of the agent 28 (or group of agents) the CPU 22 instructs the switch 26 to internally connect the port of the incoming trunk to a port of one of the identified agents.

Where the call has been connected to an agent, the CPU 22 stores the port number of the identified agent

in the termination table for the port of the incoming trunk. Likewise, the CPU 22 stores the port identifier of the incoming trunk in the termination table of the identified agent.

To complete set-up of the call to the identified agent, the CPU 22 sends a call completion message to the host 16. The call completion message may include a header identifying the message as being a call completion message. The call completion message may also include at least a port identifier of the identified agent and the call identifier. The information of the call completion message is stored in the call record previously created in conjunction with call arrival. The port identifier and call identifier allows the host 16 to deliver customer data to the specific display terminal of the agent to which the call was delivered.

Included within the host 16 is one or more NNs 32 which control system resources (e.g., the number of agents in a particular agent group, the number of incoming calls per time period accepted by the system 10, the number of outgoing calls initiated by the system 10, the number of calls in the queue of each group, etc.). Each NN 32 may be created and organized in a manner well known in the art (e.g., see "AS/400 Neural Network Utility: User's Guide and Reference PRPQ P84189").

The NN 32 may perform an analysis of resource allocation periodically (e.g., one every 10 seconds, once an hour, etc.), or upon the receipt of each call request. As used herein, a call request may be an

incoming call from the PSTN 14 or an outgoing call initiated by the host 16 as part of a call campaign.

Under the illustrated embodiment, the NN 32 may receive input information on system operation from any number of sources (e.g., from the CD 20, supervisors station 34, etc.). Information received from the CD 20 may include: the number of agents signed onto the system 10, the number of agents available for accepting calls, the delay time which a call experiences before delivery to an agent, an average time a call experiences before answering, an average talk time, a time of day, day of week, etc.

Based upon that information, the NN 32 may assign agents to groups or transfer agents among groups, based upon group loading. The NN 32 may also initiate and control outgoing call campaigns in conjunction with the control of the agent groups. Further, the NN 32 may also control the rate at which the CD 20 accepts calls from the PSTN 14.

FIG. 2 is a flow chart of neural network operation. As shown, the NN 32 may begin by processing 100 the input information to determine a system status. Following a determination of system status, the NN 32 checks to see if group resources are adequate 102.

The determination of whether or not group resources are adequate may be determined by examination of a call queue for each group. Where calls remain in the call queue for an excessive period of time, the NN 32 may determine that the resources of the group are inadequate and that more resources should be allocated to the group.

For example, the NN 32 may first identify the group with the longest and shortest delay of calls in its respective call queue. If the difference is one second or less the NN 32 may determine that the resources are adequate and that no action is needed. Alternatively, if the delay is larger (e.g., 10 seconds), the NN 32 may determine that the group with the longest delay has inadequate resources and may re-assign 104 one or more agents of the group with the shortest delay to the group with the longest delay.

Upon considering the resources of each group, the NN 32 may next consider whether or not to process a call. Processing a call may mean accepting an incoming call from the PSTN 14 or initiating an outgoing call. While the processing of incoming calls may be given priority over outgoing calls, it should be understood that either type may be given priority.

The NN 32 may determine the appropriateness of processing another call by again examining such things as a length of time a call spends in a call queue. If the time spent is determined to be excessive, the NN 32 may defer processing of another call. If the time spent in the queue is not excessive, the NN 32 may determine that another call should be processed.

Where incoming calls are given priority, the NN 32 may first check to see if there is a call waiting to be answered. If there is a call waiting to be answered, the NN 32 may signal the CD 32 to answer the call.

If there isn't a call waiting to be answered, the NN 32 may consider whether or not to initiate an outgoing call. In determining whether to initiate an outgoing call, a different set of weights may be used

for reaching the determination. For example, it may be determined that outgoing calls should not be initiated before 8am or after 5pm or during the weekend. Consequently, the inputs associated with time of day and day of week would have a different neural network weight depending on the type of call processed.

Alternatively, the NN 32 may be used to control a ratio of incoming to outgoing calls. For instance, where an operator of the system 10 is a marketing organization, it may be more important (and profitable) to ensure that more calls go out than come in. Under this model, a NN 32 may be used to first initiate calls with some ratio (e.g., 2:1, 3:1, etc.) of outgoing to incoming calls.

FIG. 3 is an example of a NN 32 that may be used for a determination of agent resources. Such a NN 32 may be maintained for each agent group associated with the ACD 10. As shown, the NN 32 for agent resources may periodically process n inputs to determine an agent resource status. Where it is determined that the group does not have enough agents, an "ADD AGENT" output may go to a "1" state to indicate the need to add an agent to a group. Where it is determined that there are too many agents in a group, a "DELETE AGENT" may go to a "1" state. The weights associated with adding agents may be different than the weights associated with deleting agents.

Where the NN 32 of one agent group signals a need for agents, the host 16 may simply looks for another group which is indicating that it has too many agents. Where a match is found, the host 16 may compare qualifications of the agents of the two groups and

transfer an agent with the appropriate qualification. Once such a determination has been made, the host 16 may send a message to the CD 20 identifying the agent and indicating the source and destination agent groups.

FIG. 4 depicts a NN 32 that may be used for call processing. As shown, a first set of information inputs 1-n may be used for a determination of call acceptance. A determination to accept a call would cause the output "CALL ACCEPT" to go high.

Also shown, is an "INITIATE CALL" output. While the same set of inputs is shown as being used for "INITIATE CALL" and "CALL ACCEPT", it is to be understood that some inputs may used for "INITIATE CALL" and not used for "CALL ACCEPT" and visa versa. This may be accomplished by having a weighting value of zero in the one case and not the other, and visa versa.

The NNs 32 may be created in a manner well-known to those of skill in the art. For example, the host 16 may be an AS/400 by IBM with a software library containing a neural network utility. An Application Programming Interface (API) may be provided for the benefit of a supervisor working at a supervisors station 34 which allows the supervisor to access and use the neural network utility program. Prompts may be provided to enable the supervisor to select a neural network model such as the back propagation model.

The programmer may also be provided with prompts to
initialize and train the neural network. Following initialization, the user may be prompted to identify a database within the host 16, which may be used to train each NN 32. Other screens may be provided through

which the supervisor may input information and adjust the weights given to the input information.

Once the NN 32 has been trained, the NN 32 may be used to allocate system resources, as appropriate. The supervisor, working through the supervisors station 34 may periodically monitor and adjust weighting values of the NN 32 as appropriate to system performance.

A specific embodiment of a method and apparatus of a neural network for controlling calls in a telephone switch according to the present invention has been described for the purpose of illustrating the manner in which the invention is made and used. It should be understood that the implementation of other variations and modifications of the invention and its various aspects will be apparent to one skilled in the art, and that the invention is not limited by the specific embodiments described. Therefore, it is contemplated to cover the present invention any and all modifications, variations, or equivalents that fall within the true spirit and scope of the basic underlying principles disclosed and claimed herein.